LA-UR -81-1640

CONF-810523--6

TITLE: AN APPLICATION OF PARALLEL PREPROCESSORS IN DATA ACQUISITION

MASTER

AUTHOR(S): H. S. Butler, M. D. Cooper, R. A. Williams,
E. B. Hughes, J. R. Rolfe, S. L. Wilson, and H. D. Zeman

SUBMITTED TO: Topical Conference on Computerized Data Acquisition
in Particle and Nuclear Physics
Oak Ridge National Laboratory, Oak Ridge, TN
May 28-30, 1981

University of California

LOS ALAMOS SCIENTIFIC LABORATORY
Post Office Box 1663 Los Alamos, New Mexico 87545
An Affirmative Action/Equal Opportunity Employer

Form No. 836 R3
St. No. 2629
12/78

UNITED STATES
DEPARTMENT OF ENERGY
CONTRACT W-7405-ENG. 36

AN APPLICATION OF PARALLEL PREPROCESSORS IN DATA ACQUISITION

H. S. Butler, M. D. Cooper, and R. A. Williams
Los Alamos National Laboratory*
Los Alamos, New Mexico 87545

E. B. Hughes, J. R. Rolfe, S. L. Wilson, and H. D. Zeman
W. W. Hansen High Energy Physics Laboratory†
Stanford University
Stanford, California 94305

Summary

A data-acquisition system is being developed for a large-scale experiment at LAMPF. It will make use of four microprocessors running in parallel to acquire and preprocess data from 432 photomultiplier tubes (PMT) attached to 396 NaI crystals. The microprocessors are LSI-11/23s operating through CAMAC Auxiliary Crate Controllers (ACC). Data acquired by the microprocessors will be collected through a programmable Branch Driver (MBD) which also will read data from 52 scintillators (88 PMTs) and 728 wires comprising a drift chamber. The MBD will transfer data from each event into a PDP-11/44 for further processing and taping. The microprocessors will perform the secondary function of monitoring the calibration of the NaI PMTs. A special trigger circuit allows the system to stack data from a second event while the first is still being processed. Major components of the system were tested in April 1981. Timing measurements from this test are reported.

Introduction

A major experiment is being mounted at the Clinton P. Anderson Meson Physics Facility (LAMPF) to investigate extremely rare processes in particle physics.[1,2] The experimental apparatus is shown in Fig. 1.

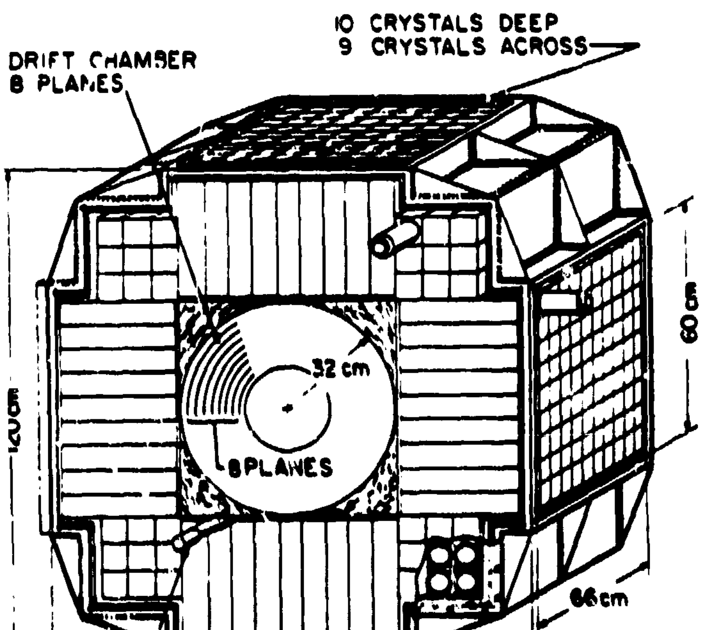

Fig. 1. Conceptual drawing of the Crystal Box detector to be used to search for several rare decay modes of the muon.

It consists of a large solid-angle modular sodium iodide detector surrounding a wire chamber and trigger hodoscope. The so-called Crystal Box will be used to search for lepton flavor-changing decays of the positive muon to $e^+\gamma$, $e^+\gamma\gamma$, and $e^+e^+e^-$ with a sensitivity to branching ratios in the range of 10^{-11}. This paper describes the data-acquisition system being assembled and tested to support the experimental program.

Data Rate

The data-acquisition system for any given experiment can be characterized by the data rate--the product of the event rate and the number of signals per event. For the Crystal Box experiment, the event rate is based on the need to sample 2×10^{12} muon decays. This total number can be achieved by balancing the running time of the experiment against the average muon stopping rate. While running at the highest possible muon stopping rate leads to the shortest experiment, an analysis[3] of the raw trigger rates from the three types of events under study showed that a number of considerations lead to an optimum muon stopping rate of 5×10^5 μ/s.** This rate yields a raw trigger rate of about 240 Hz. The corresponding data-taking time for the experiment is 1000 hours.

The signal list for the Crystal Box contains nearly 2000 signals, 90% of which are analog in origin. Each event generates pulse amplitude and time information from 432 sodium iodide channels, pulse amplitude and time information from 88 scintillator channels, time information from 728 drift chamber wires, and assorted latch bits and status words. Thus, a single event could produce about 1800 16-bit words of data. Additional information is available from 120 scaler channels, 16 temperature channels, and 28 magnets in the secondary beam line. This latter information will be recorded on a time scale measured in minutes or longer.

Attenuated Data Rate

Given the predicted event rate of 240 Hz and the signal list of 1800 words/event, the raw data rate is seen to be 4×10^5 words/s. This average rate is an order of magnitude greater than the recording rate of moderately priced magnetic tape drives, the traditional medium for recording experimental data. Moreover, during the 1000 hours of running time (4×10^6 s), the experiment would accumulate 90,000 tapes.# Experience

* Work supported by the U. S. Department of Energy

† Work supported by the National Science Foundation and the U. S. Department of Energy

**The LAMPF accelerator pulses 120 times/second for 500-700 μs giving a duty factor of 6-9%. The time between pulses is about 8 ms.

A 2400-ft reel of tape recorded at 1600 bpi and blocked in 2000-word records to give a packing fraction of 77% will hold 1.8×10^7 16-bit words.

from previous experiments at LAMPF has shown that 500 magnetic tapes is a reasonable limit on the amount of data that can be analyzed with the resources and manpower available. Thus, some combination of hardware and software preprocessing is required to attentuate the average raw event rate by a factor of 200 or better.

The first factor of 10 is obtained from a fast trigger processor. This trigger hardware imposes voter coincidence logic, stringent timing criteria, and a non-adjacency requirement on the trigger scintillators. Sophisticated geometric requirements are imposed by PROM-based matrix logic. Energy and vector momentum cuts are applied by a hardwired processor using LSI flash ADCs and digital arithmetic logic. The resulting system is described elsewhere.[3]

Another factor of 4 is obtained by data compaction. Of the 1800 words of data generated by each event, about 75% of them will have the value "zero" reflecting the fact that the channels are empty. In the drift chamber, for example, a consideration of the kinematics indicates that an average of about 70 non-zero signals will be generated for each event. To take advantage of this fact, a search was instigated for a CAMAC-based system to perform a sparse data scan on the drift chamber wires. The LRS-4290 system from LeCroy appeared to meet our requirements and on the basis of preliminary tests of a prototype model, we purchased a system.

In the case of the sodium iodide crystals, calculations and preliminary tests indicate that only about 100 crystals will be hit, yielding 100 words of pulse amplitude data and the corresponding 100 words of time data. Thus, some form of sparse data scan could be used to compress the crystal data. The exact implementation of a scheme was left to the system designers.

When the volume of compressed data from the crystals (200 words) and drift chamber (70 words) was added to the 160 words of pulse height and time information from the scintillators, the amount of data per event totalled 430 words. Various status, tag, and other identifying information pushed this total to about 450 words/event, down a factor of 4 from the initial 1800 words/event.

Design Criteria

The preceding discussion of data rates and volume of data led to a set of general criteria on the design of the data-acquisition system.

1) The system must have capabilities to acquire, preprocess, histogram, record, display, and analyze the data.

2) The system must provide for monitoring and exercising the apparatus and for diagnosing faults in it.

3) The system must reduce the recorded event rate by another factor of 5, preferably 10 (10 x 4 x 5 = 200).

4) Finally, the system must perform a sparse data scan on the crystal data.

Design Considerations

Because the final design of a system is often of less interest than the reasons for the design, a major portion of this paper is devoted to a review of the considerations which led to our design.

Systems Approach

An experiment with the number and diversity of signals of the Crystal Box requires a systems approach to the design of the data-acquisition system. The subsystems which acquire data from the major parts of the experiment--crystals, drift chamber, and scintillators--must be integrated and synchronized into a system which moves the data smoothly from the apparatus to a magnetic tape. If any subsystem fails to work in concert with all the others, then the system as a whole will stall.

Boundary Conditions

From a systems point of view, the design for the Crystal Box data-acquisition system started with a significant boundary condition, namely, the existence of a LAMPF standard data-acquisition system [4,5,6,7] which has been developed over the past 11 years. Since the standard system was designed to satisfy criteria 1) and 2) above, there was a significant motivation to adopt it for the Crystal Box experiment. By fitting our special requirements within the framework of the standard, we could minimize the amount of effort, especially programming, that would be required to get the experiment on the air and to support it during the course of two or more years of running.

Accepting the LAMPF standard meant that the computer would be a PDP-11; the readout system would be CAMAC-based, and the branch driver would be a Bi Ra Microprogrammed Branch Driver (MBD).[8,9] The programming system system had to be compatible with Program Q,[10,11] a general-purpose data-acquisition and analysis program which provides most of the functionality needed to support an experiment. Implicit in the acceptance of Program Q was the use of the RSX-11D operating system because many features of the Q program modules were optimized to run under RSX-11D. (Q is being converted[12] to run under RSX-11M; consequently, much of the programming for the Crystal Box is being done with the idea of converting to RSX-11M before the experiment is completed.)

The configuration for the PDP-11 system is given in Table I. With three exceptions, it is the defacto LAMPF standard configuration. These exceptions are: 1) a large volume disk justified by the volume of the programming system and the histogram arrays, 2) a second display scope warranted by the variety and number of graphic displays of the Crystal Box, and 3) a DECNET link to facilitate on-line data storage, analysis, and program development. The PDP-11/44 was specified to provide the required performance and to allow for more than 128K words of memory. This computer is scheduled for delivery in the Fall of 1981.

TABLE I

COMPUTER CONFIGURATION FOR THE CRYSTAL BOX EXPERIMENT

1. PDP-11 CPU with all options required by RSX-11D
2. 128K words of memory
3. Floating point unit
4. Real-time calendar clock
5. DECwriter
6. Two RK05 disks and controller
7. Large-volume disk (20M bytes) and controller
8. Two 1600-bpi, 9-track, DEC-compatible tape drives and controller
9. Two Tektronix 4010 graphic display scopes
10. Printer/Plotter, 300 lpm
11. DECNET link to VAX
12. MBD and LAMPF Trigger Module

Timing Tests

The adoption of the LAMPF standard system satisfied design criteria 1) and 2) and brought us to a consideration of 3)--the requirement to reduce the recorded event rate by a factor of 5 or better. This reduction will be achieved by applying a series of cuts to the raw data. It was tempting to off-load these calculations to the MBD, the first place where all of the data is available. However, the proposed cuts made extensive use of floating point operations and the MBD does not even have a hardware multiply/divide instruction. Hence, this load had to be assigned to the PDP-11.

To understand the implication of this load, a representative cut called the z-axis geometry cut was programmed and executed on several different models of the PDP-11. When the measured execution time was multiplied by a suitable factor to account for a total of five cuts with each event being processed through half of them, it was learned that at an even rate of 10-20/s, the time-critical processing would occupy 15-30% of a PDP-11/45 CPU. If the load exceeds this figure, then there will not be enough CPU capacity left to handle the myriad of non-time-critical but essential tasks such as displays, detailed analyses, histogramming, diagnostics, monitoring, software development, etc.

Two conclusions were derived from this timing test: 1) The PDP-11 could perform the cuts to satisfy criteria 3), but a PDP-11/45 CPU or better was needed to handle the load. Cache memory was recommended as a cost-effective way to improve the performance of the CPU by 25-40% for a margin of safety. 2) A single CPU could not perform both the cuts and the sparse data scan even if a PDP-11/70 mainframe were assumed.

Sparse Data Scan

The MBD was considered as a device to accomplish the sparse data scan and thereby satisfy criteria 4). However, timing estimates showed that while the MBD could collect the 450 words of compressed data in times less than the 8 ms between beam pulses, it could not, in addition, read out all of the crystal data and perform the sparse data scan. This conclusion suggested that additional parallel intelligence was needed beyond the MBD--the parallelism for speed and the intelligence for the sparse data scan.

Choice of Intelligence

Since CAMAC was the LAMPF standard for interfacing, the only logical choice for distributing intelligence was the Auxiliary Crate Controller (ACC). Similarly, since the LAMPF standard computer was a PDP-11, the natural choice for local intelligence was the LSI-11 because of the extensive facilities available for program development. The functions which this hardware had to perform and which dictated the final selection of commercial equipment are listed below.

1) Block Transfer: Blocks of words have to be transferred from the PDP-11 via the MBD to the LSI-11 and vice versa. Programs and tables of thresholds have to be downloaded into the LSI-11; crystal data which survived the sparse data scan have to be transferred to the PDP-11.

2) Control Functions: The intelligent device has to be able to set/read bits in a register in a proper CAMAC module which the MBD can also read and write via FCNA commands. This capability is required for the MBD and LSI-11 to communicate and synchronize their actions.

The control functions are available in several ACCs produced commercially. The block transfer capability via the DMA channel was developed at LAMPF[13] and was available commercially from only one company at the time. Thus, the modules chosen were a Bi Ra 1150 Control Port (ACC) and a Bi Ra 1151 Dataway Access Port for the DMA block transfers.

The choice of a CPU was between the model /02 and /23 of the DEC LSI-11. Timing studies based on the quoted instruction times for the 11/02 indicated that a factor of 2-3 enhancement in speed was needed to satisfy the timing constraints discussed later. Thus, the LSI-11/23 was specified.

Selecting the Bi Ra 1150 ACC left an option for choosing the packaging of the LSI. The CPU could be inserted in a CAMAC module and located in the crate or it could reside in a separate rack-mounted chassis. The costs of either approach were comparable, but an engineering choice was made in favor of the separate chassis to save slot space in the CAMAC crate and to open a wider range of suppliers.

Crystal Readout

The readout of the data from the crystals involves an integrate-and-hold circuit for the pulse amplitude and a time-to-amplitude converter circuit for the time information in each crystal. These analog signals are multiplexed under computer control to a precision ADC (Tracor Northern TN-1213) which digitizes the signals. The digitized result is read into the LSI. When all of the signals have been read, the LSI performs the sparse data scan, tags each data word with a crystal ID code, and then signals the MBD that the data are ready for transfer. The control of the multiplexing and digitization could have been implemented through CAMAC or directly through a parallel interface in the LSI. The latter was chosen to minimize the load on the Dataway.

The interface to the ADC was designed for programming simplicity and speed. The LSI issues a "MOV ADC word to Memory." The interface holds on to the Q-Bus until the digitization is complete and then transfers the word to memory. This scheme avoids interrupts or busy flags to signal that the digitization is complete and is feasible because the LSI is purpose-dedicated during the data-acquisition phase.

Final Configuration

The final system configuration is shown in Fig. 2. While only one-fourth of the equipment is pictured, there is, in fact, an intelligent multiplexer for each quadrant of the Crystal Box. The reason for parallelism and the justification for the added equipment was an execution time estimate that indicated the readout, digitization, and processing would take 3-5 ms/quadrant. Since we are trying to read out all quadrants in the 8 ms between beam pulses, the parallelism was mandatory. The readout from the ACC to the MBD is expected to take about 1 ms/quadrant.

Stabilization

The existence of the intelligent multiplexers provided an easy way to monitor and compensate for drifts in the electronics which affect the energy calibration through the photomultiplier tube (PMT) gains. The process of stabilization calls for a periodic comparing of the singles spectrum for each crystal with the corresponding reference spectrum collected and recorded at the time the PMT gain was calibrated. From this comparison, the change in gain can be calculated and folded back into the thresholds to compensate for any drift.

The stabilization spectra for the gains and pedestals will be accumulated as histograms of singles events. A separate stabilization trigger is provided to distinguish a singles event from a rare decay trigger.* Preference is given to the latter by restricting acceptance of a singles event to the last 5% of those pulses during which no rare decay occurred and no processing of data from a previous rare decay is in progress.

The memory capacity requ[illegible] [illegible]r the histograms is non-trivial. The gain spect[illegible] require 512 bins (words) each and the pedestals take another 64 bins each. For 108 crystals per quadrant, these figures lead to a total of 60K words of histogram storage in each LSI-11.

Timing studies indicate that the LSI can acquire and histogram the stabilization data in less than the 8 ms between beam pulses. Such would not be the situation unless we had one LSI per quadrant. When the spectra have enough counts, the data-taking will be turned off and the histograms will be read into the PDP-11. After the gains have been calculated, the new reference spectra stored, and the gain changes downloaded into each LSI, the system will be ready again to acquire data.

Dead Time

One final facet of the design warrants comment. Because of the length of the experiment, dead time in the apparatus had to be kept to a minimum. When the times for collecting all of the data from a single event were added up, it was determined that the data readout would take longer than the 8 ms between pulses.

The consequences of this possibility were studied in a Monte Carlo calculation which modeled the arrival of events vs. the time it took to dispose of them. When there was no buffering, i.e., the system was dead until the last event was fully processed, the dead time was found to be in the range of 13-23% for an average arrival rate of 18/s and a disposal time of 9-24 ms. This loss of beam time was judged excessive.

When a level of buffering was added, i.e., the equipment could acquire and hold the data from a second event while data from the first were still being processed, the dead time dropped an order of magnitude to, for example, less than 1% for 18/s when 16 ms were required to dispose of the data. This finding led us to devise a trigger logic that made it possible to capture the data from two, but not necessarily three, consecutive pulses.

System Operation

In late April 1981 a test system was configured to mock up the major aspects of the final system. Only one quadrant was instrumented (actually only 49 crystals) because the full complement of crystals has not been delivered. Similarly, the drift chamber was not available and that part of the experiment had to be simulated. Figure 3 shows the essential elements of the test system whose operation is described below.

The operation of the system begins in a quiescent state with no data to be collected and no processing in progress. In that state, the "Crystal Data-Acquisition System Ready" signal will be True. If, in addition, the "Computer Ready" is True (signifying that all program loading, initialization, and testing was successfully completed) and the "Operator Ready" is True (indicating that the experimenter has verified that the equipment is ready to run), then the data-acquisition system is enabled to take data.

Two types of triggers are recognized--a rare decay trigger signalled by the Master Trigger and a stabilization trigger signalled separately. When a Master Trigger occurs during the first 95% of the beam pulse, an Event Trigger is generated. It causes the ADCs and TDCs to capture the scintillator data, starts the LRS-4299 on its cycle of digitizing, sparse data scan, and acquisition, and initiates the acquisition of data in the LSI. The Event Trigger also goes to the LAMPF Trigger Module where it performs two functions. First, it generates a Busy signal on Q-Event 9, thereby disabling the trigger logic from taking further events. Second, it initiates execution in priority order of certain programs residing in the MBD. In this example, the program for Q-Event 10 runs first to collect the scintillator data. At the end of this program, the scintillator ADCs and TDCs are reenabled for the next event.

Following Q-Event 10, the program for Q-Event 9 goes into execution. It tests the LRS-4299 to determine if that module has completed the compaction of the drift chamber data. When the data are ready, Q-Event 9 reads it into the MBD. At the conclusion of the acquisition process, the drift chamber system is reenabled for the next pulse. The last action Q-Event 9 takes is to trigger Q-Events 12 and 11 internally under progr[illegible] control.

As soon as Q-Event 9 terminates, the signal "Event 9 Busy" goes False and enables one leg of the associated AND gate. The second leg of the gate is enabled when the LSI signals that the crystal data have been read by setting the "Crystal Data Collected" flag bit in a binary output register that generates a NIM signal, e.g., a Jorway-41 module. The AND of the two signals reenables the acquisition system for the next event even though all of the data from the first event have not yet been collected into the MBD.

At the time the Q-Event 12 program goes into execution, the crystal data from the event probably have not been processed in the LSI. The MBD tests a bit in the binary output module (or in the LSI itself) until the bit is set True to indicate the data have been processed. The MBD reads the data into its memory and resets the flag to signify to the LSI that the data have been retrieved. As soon as Q-Event 12 is finished, Q-Event 11 dumps the MBD buffer containing all the data for the physical event into the PDP-11 for further processing.

At the end of this cycle (Q-Events 10, 9, 12, and 11), the system is again quiescent provided no event occurred during the succeeding pulse. If, on the other hand, a second rare decay was triggered while the LSI was still processing the first, the LSI would interrupt its processing, acquire the data for the second event, and then return to processing the first one. During this time, Q-Event 12 would be waiting for the processed data from the first event. Since it has a higher priority than Q-Events 10 and 9. the acquisition system will not be reenabled until the first event is read into the MBD. This is true because Q-Event 9 which was triggered by the second event went into a Busy state which witholds the enable from the system. Thus, the system can acquire data from two, but not necessarily three, consecutive pulses. At an average event rate of 10-20/s, the probability of a Master Trigger occurring in three consecutive pulses is acceptably low so that the dead time is negligibly small.

The final operation of the system is stabilization. If a rare decay does not occur in the first 95% of the pulse, then the Event Trigger is disabled and a stabilization trigger is enabled unless the LSI is still

*The expression "rare decay" is used in the rest of the paper to mean an event which is potentially a rare decay, i.e., one which has passed all the tests to this point as signalled by a Master Trigger.

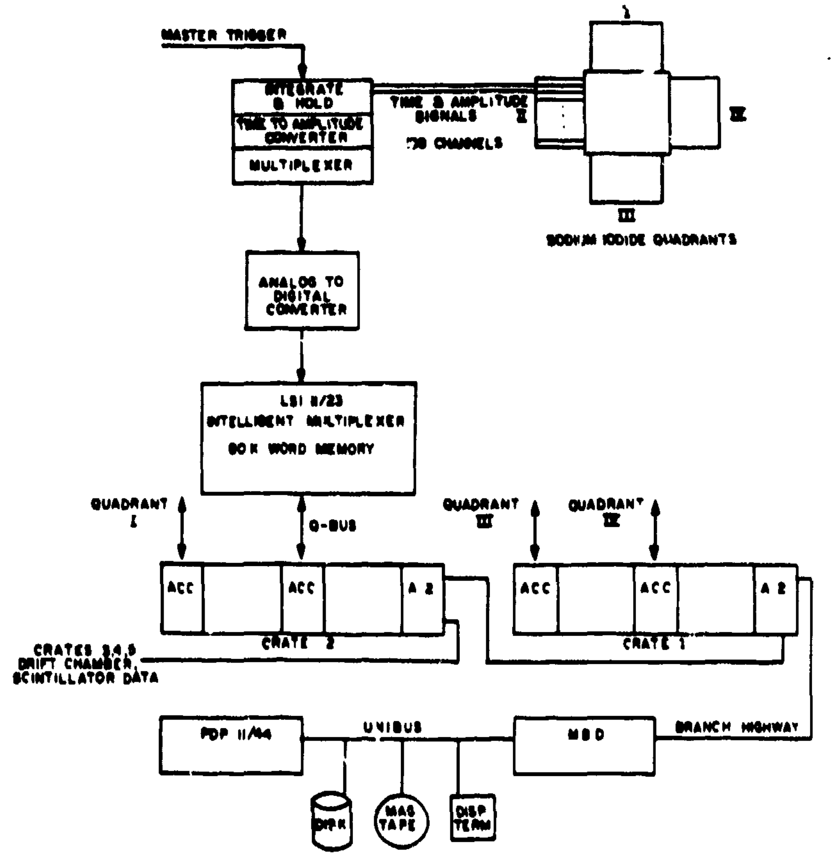

Fig. 2. Schematic representation of the Crystal Box data-acquisition system. Signals from the sodium iodide crystals in each quadrant are acquired through the Intelligent multiplexers. Data from other portions of the experiment are acquired through the CAMAC Branch Highway. The MBD combines all portions of the data for transfer to the PDP-11/44.

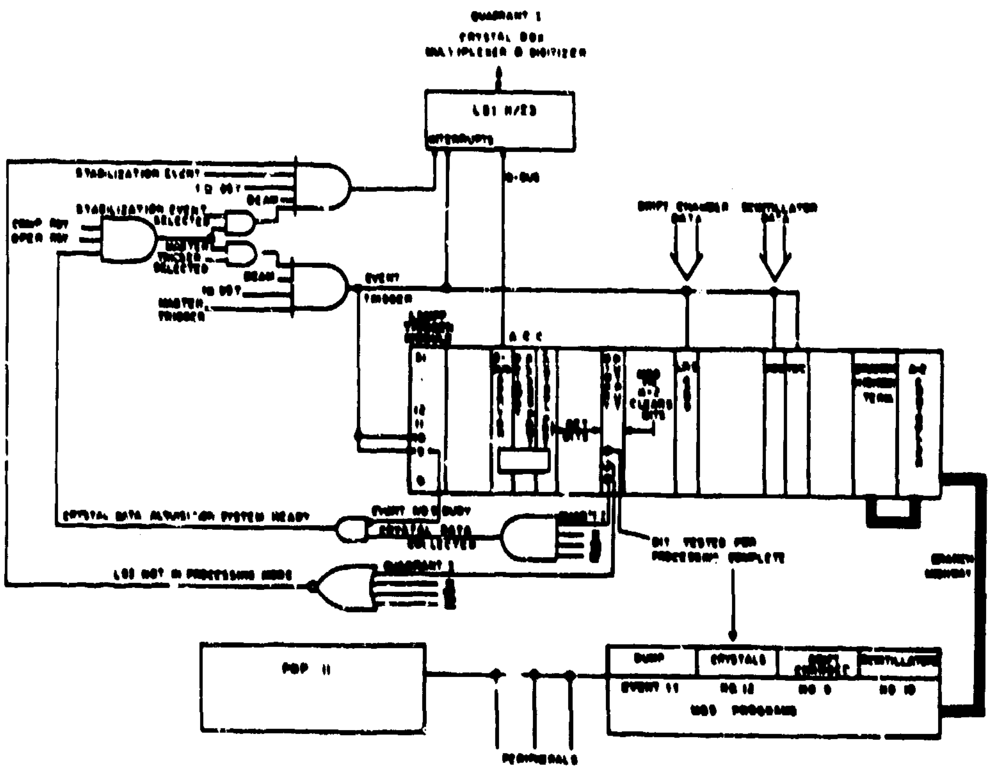

Fig. 3. Schematic diagram of the test system which implemented the important aspects of the proposed system. The LAMPF Trigger Module synchronizes the operation of the system. The logic external to the CAMAC crate allows the system to capture data from a second event while the first is still being processed.

processing data from the previous decay as signalled by a bit set in the binary output module. It is required that a stabilization event be acquired and histogrammed in less than the 8 ms between beam pulses.

The extension from one quadrant pictured here to four quadrants is simply to replace Q-Event 12 with a sequence of Q-Events 15, 14, 13, and 12, each reading data from a separate LSI.

Programming System

The programming system in the PDP-11 is an integral part of the data-acquisition system it supports. The flow of data within the PDP-11 is shown in Fig. 4.

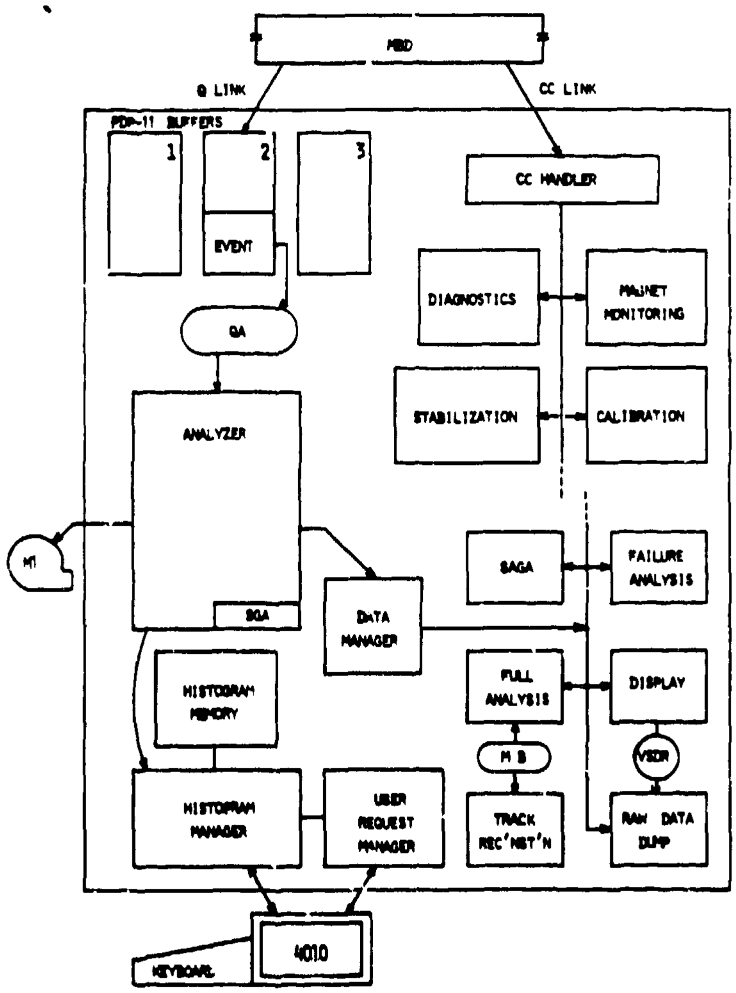

Fig. 4. Schematic representation of the programming system for the PDP-11. Event data is passed from the MBD to the Analyzer for processing, histogramming, and taping. Secondary processing is handled by programs in the bottom right quadrant. They obtain event data via the Data Manager. The CC Handler is used for control and acquisition of data other than event data.

The so-called Dump Event, 11, causes the data in the MBD to be transferred to the PDP-11 via a DMA transfer. This block of data goes into one of three buffers within the PDP-11. At the appropriate time, a module of the Q program called QA distributes the data for one physical event to the Analyzer, a PDP-11 program running at a moderately high priority level. As a first step, the Analyzer performs certain validity checks on the data to be sure the equipment is working and that Q-Events from two consecutive physical events did not get mixed. Next, the Analyzer will apply various cuts to the data and, if the event passes these cuts, the data will be recorded on magnetic tape. On the basis of other cuts, various segments of the data will be histogrammed. This process is handled through the Q Histogram Manager[14] which is set up at the start of the run from tables associated with the User Request Manager.

In addition to the immediate processing performed by the Analyzer, there is subsequent processing by programs operating at a lower priority. These programs process data acquired from the Analyzer. To effect this transfer, a mechanism had to be found to pass data from a higher priority program (the Analyzer) to a lower priority program, e.g., a display program. The mechanism adopted was one developed by Wayne Kinnison.[15] The mechanism sets up several data areas within the Analyzer during initialization of the system. When a program requires data from one of these areas, it issues a request via the RSX Executive to the Data Manager and then suspends execution. The Data Manager has to wait until the Analyzer is not running before it can get control. Its first action is to disable the Analyzer to be sure no interruptions occur during the transfer. Then it transfers the data, issues a Resume to the requesting program, and reenables the Analyzer. This mechanism ensures that data from successive physical events do not get mixed. It also minimizes any conflict with the Analyzer, i.e., at worst, the Analyzer is delayed briefly while the Data Manager is transferring data.

The programs in the lower right quadrant of the drawing are typical of the secondary processing that can take place. They are activated by the experimenter as required. If they have to pass data among themselves, they use the mechanisms of RSX-11D, e.g., VSDR, or a Mail Box handler developed at LAMPF for passing large blocks of data in a non-interferring way.

Finally, there is a class of tasks in the upper right quadrant which makes use of a special CAMAC handler, CC, to issue control instructions or acquire data via non-Q-related codes in the MBD. For example, the stabilization histograms from the LSI are acquired through this port rather than the Q Link. These operations are generally initiated by the experimenter rather than by some trigger.

Test Results

During the run in April, we had the opportunity to test certain facets of the system shown in Fig. 3, mainly those dealing with the LSI and the ACC. After some initial difficulties, we were able to acquire data from a collection of 49 crystals along with data from scintillators and other devices. Of particular interest were measurements of the time to handle stabilization events and to read out the crystal data. These times indicate that stabilization events can be handled in less than 8 ms as required and that acquisition times for rare decays can be kept under 16 ms (2 beam pulses). By making the readout of the ADC more efficient and by optimizing the coding in the LSI, it appears entirely feasible to meet the timing requirements. Further tests are scheduled for July and August; however, equipment for the full four quadrants has been ordered already.

Acknowledgments

The authors wish to acknowledge the technical advice and assistance of Jim Amann, Rick Floyd, Jim Harrison, Tom Kozlowski, Don Machen, Howard Matis, Mike Oothoudt, and Jerry Potter of the Los Alamos National Laboratory.

References

1. "Crystal Box for Rare Decays," CERN Courier, pp 312, 313 (October 1979).

2. Cy Hoffman and Minh Duong-Van, "The Search for Muon-Number Violation at LAMPF," Los Alamos Science, pp 62-67 (Summer 1980).

3. G. H. Sanders et al., "Intelligent Trigger Processor for the Crystal Box," Proceedings of the Topical Conference on the Application of Microprocessors to High-Energy Physics Experiments, CERN, Geneva (May 4-6, 1981).

4. H. L. Anderson et al., "LAMPF Data-Acquisition System," Los Alamos Scientific Laboratory report LA-4504-MS (August 1970).

5. R. F. Thomas, Jr. et al., "Computing in Suppport of Experiments at LAMPF," Los Alamos Scientific Laboratory report LA-6543-MS (October 1976).

6. D. G. Perry et al., "Data Analysis Facility at LAMPF," Los Alamos Scientific Laboratory report LA-7034-MS (November 1977).

7. D. G. Perry, "The LAMPF Standard Data-Acquisition System," IEEE Trans. Nuc. Sci., Vol. NS-26, No. 4, pp 4422-4426 (August 1979).

8. L. R. Biswell and R. E. Rajala, "Design and Operation of a Microprogrammed Branch Driver for a PDP-11 Computer," Los Alamos Scientific Laboratory report LA-5144 (May 1973)

9. R. F. Thomas, Jr., "Programming the MBD Model II," Los Alamos Scientific Laboratory report LA-5404-MS, (September 1973).

10. Martin Kellogg et al., "Introduction to Q," Los Alamos Scientific Laboratory report LA-7001-M (June 1978).

11. D. G. Perry, "The Q System: An Experiment in Software Standardization for Data Acquisition," IEEE Trans. Nuc. Sci., Vol. NS-26, No. 4, pp 4494-4499 (August 1979).

12. James F. Harrison et al., "Design of the RSX-11M Q System," IEEE Trans. Nuc. Sci., Vol. NS-28, No. 5, (October 1981).

13. D. R. Machen, "Design and Operation of the LAMPF Auxiliary Contriller--High Speed Remote Processing on the CAMAC Dataway," Los Alamos Scientific Laboratory report LA-7683-MS (February 1979).

14. Martin Kellogg et al., "A System for Histogram Entry, Retrieval, and Plotting," Los Alamos Scientific Laboratory report LA-6920-M (October 1977).

15. Wayne Kinnison, Los Alamos National Laboratory, private communication (1980).